arXiv:1410.7171v2 [math.OC] 5 Feb 2015

Exponentiated Subgradient Algorithm for Online Optimization under the Random Permutation Model

Reza Eghbali*[1], Jon Swenson[2], and Maryam Fazel[1]

[1]Department of Electrical Engineering, University of Washington, Seattle
[2]Department of Mathematics, University of Washington, Seattle

February 6, 2015

Abstract

Online optimization problems arise in many resource allocation tasks, where the future demands for each resource and the associated utility functions change over time and are not known apriori, yet resources need to be allocated at every point in time despite the future uncertainty. In this paper, we consider online optimization problems with general concave utilities. We modify and extend an online optimization algorithm proposed by Devanur et al. for linear programming to this general setting. The model we use for the arrival of the utilities and demands is known as the random permutation model, where a fixed collection of utilities and demands are presented to the algorithm in random order. We prove that under this model the algorithm achieves a competitive ratio of $1 - O(\epsilon)$ under a near-optimal assumption that the bid to budget ratio is $O\left(\frac{\epsilon^2}{\log(m/\epsilon)}\right)$, where m is the number of resources, while enjoying a significantly lower computational cost than the optimal algorithm proposed by Kesselheim et al. We draw a connection between the proposed algorithm and subgradient methods used in convex optimization. In addition, we present numerical experiments that demonstrate the performance and speed of this algorithm in comparison to existing algorithms.

1 Introduction

Online optimization provides a framework for real-time optimization problems that arise in applications such as Internet advertising, online resource allocation, and online auctions. In these problems, functions of new variables are added to the objective, and constraints are modified over time. An online optimization algorithm should assign a value to each newly introduced variable without any knowledge of future changes in the objective function. Linear programs have received considerable attention in the context of online optimization ([AWY09], [FHK$^+$10], [DJSW11],[JL12],[KRTV14]). In this paper, we study a more general online convex optimization problem that contains *online linear programming* (online LP) as a special case. The performance of an online algorithm for any instance of the problem can be evaluated by finding the ratio P/P^*, where P is the objective value achieved by the algorithm and P^* is the optimal value of the corresponding offline optimization problem. The competitive ratio of an online maximization algorithm under the worst-case model is the infimum of P/P^* over all the possible instances. While for many special online problems algorithms have been proposed that achieve a non-trivial competitive ratio under the worst-case model, the competitive ratio of any online LP algorithm tends to zero as the number of functions grows [BIKK08]. To be able to derive non-trivial bounds on the competitive ratio of an online algorithm for any problem that contains online LP, one often considers the competitive ratio under more restrictive models such as the random permutation model or the i.i.d. model. The results of this paper are presented for the random permutation model. In section 1.1, we formulate the problem and formally define the competitive ratio of an online algorithm under various models.

1.1 Problem formulation

Consider the problem

$$\underset{x_1,\dots,x_n\in\mathbb{R}^k}{\text{maximize}} \quad \sum_{t=1}^{n} f_t(x_t) + \psi\left(\sum_{t=1}^{n} A_t x_t\right), \tag{$\mathcal{P}0$}$$

*eghbali@uw.edu

where for all $t \in [n] = \{1, 2, \ldots, n\}$, $A_t \in \mathbb{R}_+^{m \times k}$, $f_t \in \mathcal{G}$, with $\mathcal{G}$ a subset of proper concave functions mapping $\mathbb{R}^k$ to $[-\infty, +\infty)$. More technical assumptions on the functions in $\mathcal{G}$ will be specified in section 4. Here $\psi(u) = I_{u \leq b}$ for some $b \in \mathbb{R}^m$, where I denotes the indicator function of a set, defined as follows for any $U \subset \mathbb{R}^m$:

$$I_{u \in U} = \begin{cases} 0 & u \in U, \\ -\infty & u \notin U. \end{cases}$$

Often in applications, the functions f_t are utility functions, i.e., a decision x_t provides the utility $f_t(x_t)$, and the function $\psi = I_{u \leq b}$ can be interpreted as imposing a total *budget* given by $b \in \mathbb{R}^m$. In an online optimization problem, for all $t \in [n]$, a vector should be assigned to x_t that can depend on all the information obtained until time t, i.e., (f_s, A_s), $1 \leq s \leq t$, but not on the future information (f_s, A_s), $t+1 \leq s \leq n$. A procedure that performs this task is called an *online optimization algorithm*. The input to the algorithm, $((f_1, A_1), (f_2, A_2), \ldots, (f_n, A_n))$, is a random vector taking values in $\mathcal{H}$, where $\mathcal{H} \subset \left(\mathcal{G} \times \mathbb{R}_+^{m \times k}\right)^n$. To avoid issues with measurability, we assume that this random vector takes values in a countable subset of $\mathcal{H}$. Let $\mathcal{F}$ be the set of probability distributions corresponding to such random vectors. To evaluate the performance of an online optimization algorithm, one can define the *competitive ratio*. Let x_t, $1 \leq t \leq n$ be the solution given by an online algorithm and define $P = \sum_{t=1}^n f_t(x_t) + \psi\left(\sum_{t=1}^n A_t x_t\right)$. Let P^* be the optimal value of $(\mathcal{P}0)$. Assume that $P^* \in (0, +\infty)$ for all the members of $\mathcal{H}$. The competitive ratio over $\mathcal{D} \subset \mathcal{F}$ is defined as:

$$C_{\mathcal{D}} = \inf \left\{ E\left[\frac{P}{P^*}\right] \middle| ((f_1, A_1), (f_2, A_2), \ldots, (f_n, A_n)) \sim g \in \mathcal{D} \right\}.$$

Three important choices for $\mathcal{D}$ considered in the literature for special cases of problem $(\mathcal{P}0)$ are as follows:

1. Worst-case: $\mathcal{D}_1 = \mathcal{F}$. In this case, $C_{\mathcal{D}_1}$ simplifies to:

$$C_{\mathcal{D}_1} = \inf \left\{ \frac{P}{P^*} \middle| ((f_1, A_1), (f_2, A_2), \ldots, (f_n, A_n)) \in \mathcal{H} \right\}.$$

   Note that if the online algorithm assigns a random value to P (the algorithm is randomized), P should be replaced by $E[P]$.

2. Random permutation: $\mathcal{D}_2$ is the subset of $\mathcal{F}$ corresponding to exchangeable random vectors, i.e., distribution functions that are invariant under permutation. Let Π be the set of all permutations over $[n]$. Let σ be a random permutation uniformly distributed on Π. The definition of $C_{\mathcal{D}_2}$ simplifies to:

$$C_{\mathcal{D}_2} = \inf \left\{ \frac{E[P]}{P^*} \middle| ((f_1, A_1), (f_2, A_2), \ldots, (f_n, A_n)) = \left(Y_{\sigma(1)}, Y_{\sigma(2)}, \ldots, Y_{\sigma(n)}\right), (Y_1, Y_2 \ldots, Y_n) \in \mathcal{H} \right\}.$$

3. i.i.d.: $\mathcal{D}_3$ contains all $g \in \mathcal{F}$ induced by random vectors $((f_1, A_1), (f_2, A_2), \ldots, (f_n, A_n))$ such that (f_t, A_t), $1 \leq t \leq n$ are independent and identically distributed.

Note that since $\mathcal{D}_3 \subset \mathcal{D}_2 \subset \mathcal{D}_1$, $C_{\mathcal{D}_3} \geq C_{\mathcal{D}_2} \geq C_{\mathcal{D}_1}$.

1.2 Related work

Most literature in online optimization has focused on a special case of $(\mathcal{P}0)$ in which for all $t \in [n]$, the domain of f_t is the simplex $\{x \in [0,1]^k \mid \mathbf{1}^\top x \leq 1\}$ and $f_t(x) = c_t^T x$ for $x \in \mathrm{dom} f_t$. Following [AWY09], we refer to this problem as the online LP problem, expressed as

$$\begin{array}{ll} \text{maximize} & \sum_{t=1}^n c_t^T x_t \\ \text{subject to} & \sum_{t=1}^n A_t x_t \leq b \\ & \mathbf{1}^\top x_t \leq 1, x_t \geq 0 \;\; t \in [n]. \end{array} \tag{1}$$

Note that $C_{\mathcal{D}_1}$ for any online LP algorithm goes to zero when n grows as shown by an example in [BIKK08], while for the same example under the random permutation model the optimal algorithm achieves a competitive ratio that tends to $1/e$ as n goes to infinity [BIKK08]. In light of this fact, much attention has been dedicated to the random

permutation and i.i.d. models for online LP. Here we review the related results for online LP. In section 1.4, we briefly review the special cases and the applications of combinatorial online optimization problems, with convex relaxations that can be expressed as online linear programs.

An important parameter that appears in the analysis of many online LP algorithms is $\gamma = \max_{t,i,j} \frac{(A_t)_{i,j}}{b_i}$ which is called the *bid-to-budget* ratio. For an online algorithm to achieve a competitive ratio of $1-\epsilon$ under random permutation or i.i.d. models, it is necessary that $\gamma = O\left(\frac{\epsilon^2}{\log m}\right)$ when $m \geq 2$ ([AWY09], [DJSW11]). Feldman et al. [FHK$^+$10] presented an algorithm with $C_{\mathcal{D}_2} \geq 1-\epsilon$ when $\gamma = O\left(\frac{\epsilon^3}{m \log(nk)}\right)$ and $\gamma' = O\left(\frac{\epsilon}{\log(nk)}\right)$, where $\gamma' := \frac{\max_{t,i} c_{t,i}}{P^*}$. Agrawal et al. [AWY09] proposed an online LP algorithm called *dynamic learning algorithm (DLA)* with $C_{\mathcal{D}_2} = 1 - O(\epsilon)$ when $\gamma = O\left(\frac{\epsilon^2}{m \log(nk/\epsilon)}\right)$. DLA solves $\log_2(1/\epsilon)$ linear programs to estimate the optimal dual variables associated with the m constraints. The primal variable is then assigned using the estimated dual variables. The dependency that exists on n in their results is introduced by a union bound over the space of dual variables. Molinaro and Ravi [MR13] improved this union bound and proved that a modified version of DLA achieves a competitive ratio of $1 - O(\epsilon)$ when $\gamma = O\left(\frac{\epsilon^2}{m^2 \log(m/\epsilon)}\right)$. Kesselheim et al. [KRTV14] proposed an algorithm that achieves a competitive ratio of $1-\epsilon$ when $\gamma = O\left(\frac{\epsilon^2}{1+\log d}\right)$. Here d is the maximum number of non-zero elements of any column of A_t for all $t \in [n]$. When $d = m$ and $m \geq 2$ their condition gives $\gamma = O\left(\frac{\epsilon^2}{\log m}\right)$, which matches the necessary condition on γ. When $d = 1$, their condition transforms to $\gamma = O\left(\epsilon^2\right)$, which also matches the necessary condition given in [Kle05]. We call this algorithm *KRTV* (dubbed after the initials of the authors' last names). At each $t \in [n]$, KRTV chooses the value of x_t by solving a linear program with the same number of linear constraints as the original problem but with tk variables. Devanur et al. [DJSW11] proposed an online algorithm that achieves a competitive ratio of $1 - O(\epsilon)$ under a model they call the adversarial stochastic model when $\max\{\gamma, \gamma'\} = O\left(\frac{\epsilon^2}{\log(m/\epsilon)}\right)$. The adversarial stochastic model is more general than the i.i.d. model. However, their analysis does not apply to the random permutation model.

After finishing this paper, we became aware of two very recent results developed independent of this work and at around the same time. Authors in [AD14] and [GM14] have analyzed similar algorithms to the one proposed in [DJSW11]. In [GM14], the authors have analyzed this algorithm for online LP and demonstrated a competitive ratio of $1 - O(\epsilon)$ when $\max\{\gamma, \gamma'\} = O\left(\frac{\epsilon^2}{\log(m/\epsilon)}\right)$. They also prove their analysis holds true for a class of more general online linear programs that may include inequalities of form $\sum_{t=1}^{n} A_t x_t \geq b$. In [AD14], authors have provided a competitive difference analysis for their algorithm applied to a class of general convex programs while showing that a specialized version of their algorithm for problem ($\mathcal{P}0$) achieves a competitive ratio of $1-\epsilon$ when $\max\{\gamma, \gamma'\} = O\left(\frac{\epsilon^2}{\log(m)}\right)$.

1.3 Our results

In this paper, we propose a modified and extended version of the online LP algorithm in [DJSW11] that applies to the more general problem ($\mathcal{P}0$). We prove that this algorithm achieves a competitive ratio of $1 - O(\epsilon)$ under the random permutation model when $\max\{\gamma, \gamma'\} = O\left(\frac{\epsilon^2}{\log(m/\epsilon)}\right)$ (the generalization of definition of γ for ($\mathcal{P}0$) is given in section 4). This condition is the same as the condition given in [DJSW11] for the adversarial stochastic model.

We interpret the proposed algorithm as a form of subgradient descent, and call it *exponentiated subgradient algorithm (ESA)*. ESA solves only $\log_2(1/\epsilon)$ optimization problems with number of variables $nk\epsilon, 2nk\epsilon, 4nk\epsilon, \ldots, nk/2$. This makes ESA computationally as efficient as DLA, while having a theoretical guarantee on its competitive ratio under the random permutation model that is close to the theoretical guarantee for KRTV.

1.4 Applications and special cases

Consider the (non-convex) problem where $f_t(x) = c_t^T x$ over a binary domain $\text{dom} f_t = \Gamma := \{x \in \{0,1\}^k \mid \mathbf{1}^\top x \leq 1\}$, which we refer to as the online allocation problem. It contains many combinatorial online problems studied in the literature such as online bipartite matching [KVV90], the AdWords problem [MSVV07], and the multiple secretary problem [Kle05]. The convex relaxation of the online allocation problem is the online LP problem. Note that all the online LP algorithms discussed above produce an integer solution; therefore, the bounds on their competitive ratio holds for the online allocation problem as well.

In many applications of online problems, the index t is associated with time. We present an example of the online allocation problem using the terminology of servers and clients. Many other online allocation problems are equivalent to this example. Consider a server that has to serve a number of clients. At time t, a client sends an order of form $(c_t, A_t) \in \mathbb{R}_+ \times \mathbb{R}_+^m$. A_t determines the amount of resources needed for the fulfillment of that order and c_t is the price that the client is willing to pay for the order. From the point of view of the server, c_t can be viewed as the utility

associated with that order. The server has to make an irrevocable decision by choosing $x_t \in \{0,1\}$. If $x_t = 1$, the order is fulfilled and the total utility is increased by c_t; otherwise, the order is rejected and the total utility does not change. The goal of the server is to maximize the total utility while not exceeding any limit on resources. In the case where $k > 1$, each column of A_t is an option. The server should make a decision of which option, if any, to fulfill by choosing $x_t \in \Gamma$. In the LP relaxation, each utility function is linear over its domain. One can assume a scenario where the order can be partially fulfilled with a utility function associated with each order that follows a diminishing return. In this case, every order is of form (f_t, A_t), where f_t, $t = 1, \ldots, n$ are concave functions. Special cases of online optimization problems where there is a relationship between c_t and A_t have been extensively studied in the literature; here we briefly review the results for four important special cases. For a thorough review of the results on these examples, we refer the reader to [Meh13].

1. Online Bipartite Matching: Consider a bipartite graph $G = ((U, V), E)$ with $|U| = n$ and $|V| = m$, where $|U|$ denotes the cardinality of U. Let $A = [a_1, a_2, \ldots, a_n] \in \{0,1\}^{m \times n}$ be the adjacency matrix for G. Online bipartite matching is a special case of online allocation problem with $k = m$, $b = \mathbf{1}$, $A_t = \sum_{i=1}^{m} a_{t,i} e_i e_i^T$, and $c_t = a_t$ for all $t \in [n]$, where $\{e_1, e_2, \ldots, e_m\}$ is the standard basis for $\mathbb{R}^m$. Karp et al. [KVV90] proposed a randomized online bipartite matching algorithm, called the ranking algorithm, that achieves $1 - 1/e$ competitive ratio under the worst-case model, which they proved is the best achievable competitive ratio by any online matching algorithm under this model. Under the random permutation model, the ranking algorithm achieves a competitive ratio of at least $0.696 > 1 - 1/e$ [MY11] and no algorithm can achieve a competitive ratio larger that 0.727 [KMT11]. Several authors have proposed online algorithms and studied their competitive ratio under the i.i.d. model when the algorithm receives the distribution as an input (see, e.g., [FMMM09], [MGS12] and [JL13]). A competitive ratio of 0.706 under this model is achieved by the algorithm proposed in [JL13].

2. AdWords: This problem is a generalization of online bipartite matching with $A \in \mathbb{R}_+^{m \times n}$ and $b \in \mathbb{R}_{++}^m$. The name of this problem comes from its application in online advertising by search engines. A realistic assumption in this problem is that γ is bounded by a small number. That renders the results for competitive ratio of online LP algorithms under i.i.d. or random permutation models applicable to this problem. In fact, the online AdWords algorithm proposed by Devanur and Hayes [DH09] for random permutation model inspired the later results for online LP algorithms. Mehta et al. provided an algorithm with $C_{\mathcal{D}_1} \to 1 - 1/e$ as $\gamma \to 0$, which is also shown to be the best competitive ratio under this model [MSVV07]. Buchbinder et al. [BJN07] proposed an algorithm for the worst-case model that explicitly utilizes the dual variables and achieves the optimal competitive ratio. This algorithm retains a vector of dual variables and chooses the primal variable x_t, according to the complementary slackness rule. The dual variables are updated after each assignment of x_t. Devanur and Jain [DJ12] proposed an algorithm for a concave version of AdWords, which linearizes the problem at each step and utilizes a similar primal-dual approach.

3. Multiple Secretary Problem: This problem has been proposed by Kleinberg [Kle05] as a variation of the classical secretary problem. It can be viewed as a special case of an online allocation problem with $k = 1$, $m = 1$, $c_t \in \mathbb{R}_+$, $A_t = 1$ and $b \in \mathbb{N}$. The pessimistic example for the competitive ratio of online LP algorithms under the worst-case model falls under this problem with $b = 1$. Therefore, this problem and many of its generalization are stated and studied under the random permutation model. Kleinberg [Kle05] proposed an algorithm with $C_{\mathcal{D}_2} = 1 - O\left(1/\sqrt{b}\right)$, which was shown to be order-wise optimal. In other words, no online algorithm can achieve $C_{\mathcal{D}_2} \geq 1 - \epsilon$ unless $\gamma = O\left(\epsilon^2\right)$. Bibaioff et al. [BIKK07] proposed a simple algorithm that achieves a competitive ratio of $\frac{1}{e}$ for any b as n tends to infinity.

4. Online Weighted Bipartite Matching: This problem is a generalization of online bipartite matching. It can also be viewed as a generalization of the multiple secretary problem. Consider a bipartite graph and its adjacency matrix $A \in \{0,1\}^{m \times n}$. Let $W = [w_1, w_2, \ldots, w_n] \in \mathbb{R}_+^{m \times n}$ be the weight matrix associated with the edges in the graph. In this problem, $k = m$, $b \in \mathbb{N}^m$, $A_t = \sum_{i=1}^{m} a_{t,i} e_i e_i^T$ and $c_t = w_t$ for all $t \in [n]$. Similar to the AdWords problem when γ is small, the results for online LP algorithms are applicable to this problem. Without any assumption on γ, $C_{\mathcal{D}_2} \leq \frac{1}{e}$ since this problem contains the secretary problem. Korula and Pál [KP09] proposed an algorithm with $C_{\mathcal{D}_2} \geq \frac{1}{8}$. Their algorithm incorporates the optimal stopping rule for the classical secretary problem. Kesselhiem et al. [KRTV13] proposed an algorithm that incorporates the same stopping rule and achieves the optimal competitive ratio of $C_{\mathcal{D}_2} = \frac{1}{e}$.

2 Preliminaries

In this section, we review basic concentration inequalities that will be used in the rest of this paper.

Proposition 1 (Bernstein's inequality). *Suppose Y_i, $1 \leq i \leq n$ are independent random variables taking values in $\mathbb{R}$. Let $Z = \sum_{i=1}^{n} Y_i$, $\mu = E[Z]$, and $v = \sum_{i=1}^{n} E[Y_i^2]$. Let $\gamma > 0$ and suppose $Y_i \leq \gamma$ for all i, then for all $r \in [0, 3/\gamma)$:*

$$\log E\left[\exp\left(r\left(Z-\mu\right)\right)\right] \leq \frac{vr^2}{2(1-\gamma r/3)};$$

therefore, for all $t > 0$:

$$\mathbb{P}\left(Z-\mu > t\right) \leq \exp\left(-\frac{t^2}{2(v+\gamma t/3)}\right).$$

For proof see theorem 2.10 and corollary 2.11 of [BLM13]. Note that if $0 \leq Y_i \leq \gamma$ for all i, then $v = \sum_{i=1}^{n} E[Y_i^2] \leq \gamma \sum_{i=1}^{n} E[Y_i] = \gamma\mu$, which simplifies the Bernstein's inequality:

$$\mathbb{P}\left(Z-\mu > t\right) \leq \exp\left(-\frac{t^2}{2\gamma(\mu+t/3)}\right), \tag{2}$$

$$\mathbb{P}\left(Z-\mu < -t\right) \leq \exp\left(-\frac{t^2}{2\gamma\mu}\right), \tag{3}$$

for all $t > 0$.

Proposition 2 (Hoeffding's theorem). *Let $X_1, \ldots, X_n$ be samples from a finite multiset $M \subset \mathbb{R}$ drawn without replacement, and let $Y_1, \ldots, Y_n$ be samples drawn with replacement from M. If f is a convex function, then:*

$$E\left[f\left(\sum_{i=1}^{n} X_i\right)\right] \leq E\left[f\left(\sum_{i=1}^{n} Y_i\right)\right]. \tag{4}$$

Using proposition 2, one can apply the result of proposition 1 to random variables sampled without replacement. Note that one can sharpen some of the bounds for the case of sampling without replacement; for example, see [Ser74].

Proposition 3 (Doob's maximal inequality). *Suppose S_i, $i = 1, 2, \ldots, n$ is a sub-martingale. Let $M_n = \max_{1 \leq i \leq n} S_i$, then for all $t > 0$:*

$$t\mathbb{P}\left(M_n > t\right) \leq E\left|S_n\right|, \tag{5}$$

$$\mathbb{P}\left(M_n > t\right) \leq \inf_{r>0} \frac{E\left[\exp\left(rS_n\right)\right]}{\exp\left(rt\right)}. \tag{6}$$

For proof see chapter 8.10 [Sho00].

In the rest of this paper, Π denotes the set of all the permutations of $[n]$, i.e., set of bijections from $[n]$ to $[n]$. For all $B \subset \Pi$, we denote the complement of B with $\bar{B}$, and we define :

$$1_B(s) = \begin{cases} 1 & s \in B, \\ 0 & s \notin B. \end{cases} \tag{7}$$

We denote the transpose of a vector v by $v^\top$. For a function $f : \mathbb{R}^k \mapsto [-\infty, +\infty)$, f^* denotes the concave conjugate of f, and is defined as:

$$f^*(v) = \inf_x v^\top x - f(x),$$

for all $v \in \mathbb{R}^k$. For a concave function f, $\partial f(x)$ denotes the set of supergradients of f at x, i.e., the set of all $v \in \mathbb{R}^k$ such that:

$$\forall x' \in \mathbb{R}^k : \quad f(x') \leq v^\top(x'-x) + f(x).$$

For a convex function f, we use the same notation $\partial f(x)$ to denote the set of subgradients of f at x, i.e., the set of all $v \in \mathbb{R}^k$ such that:

$$\forall x' \in \mathbb{R}^k : \quad f(x') \geq v^\top(x'-x) + f(x).$$

3 A feasibility problem

We start by considering a simple form of ESA applied to a linear feasibility problem to illustrate the proof technique and ideas. Let $\Delta = \{x \in \mathbb{R}^k | \mathbf{1}^\top x = 1,\ x \geq 0\}$. Consider a spacial case of $(\mathcal{P}0)$, where $b = \mathbf{1} \in \mathbb{R}^m$, and for all $t \in [n]$, $f_t(x) = I_{x \in \Delta}$. This problem can be represented as the following feasibility problem:

$$\text{Find} \quad (x_1, x_2, \ldots, x_n) \in \Delta^n \quad \text{such that} \quad \sum_{t=1}^n A_t x_t \leq \mathbf{1}, \tag{8}$$

where for all t, $A_t \in \mathbb{R}_+^{m \times k}$. Let $(x_1^*, x_2^*, \ldots, x_n^*)$ be a feasible solution for (8). Define $\gamma = \max_{t,i,j} A_{t,i,j}$ and choose σ uniformly at random from Π. Consider Algorithm 1 for problem (8).

Algorithm 1

Require: $\epsilon \in (0,1)$, $\gamma > 0$
  $\nu = \frac{1}{\gamma} \log(1+\epsilon)$
  for $t = 1 : n$ **do**
    $x_t \in \operatorname{argmin}_{z \in \Delta} \sum_{i=1}^m \exp\left(\nu \sum_{s=1}^{t-1} \left(A_{\sigma(s)} x_s\right)_i\right) \left(A_{\sigma(t)} z\right)_i$
  end for

Algorithm 1 is proposed by Devanur et al. [DJSW11]. Here we extend the analysis of algorithm 1 to the random permutation model. We first derive concentration inequalities for the running sums of form $\sum_{s=t}^n \left(A_{\sigma(s)} x^*_{\sigma(s)}\right)_i$. These inequalities will be used in the analysis of the algorithm, where we compare $(x_1, x_2, \ldots, x_n)$, the solution given by the algorithm, with $(x_1^*, x_2^*, \ldots, x_n^*)$. Let $L = \log_2(\frac{1}{\epsilon})$. To simplify the notation, we assume that L and $n\epsilon$ are integers in the rest of this paper. Let $\kappa : [n(1-\epsilon)] \mapsto [L]$ be defined as $\kappa(t) = \left\lfloor \log_2\left(\frac{n-t}{n\epsilon}\right)\right\rfloor + 1$. Equivalently, $\kappa(t) = k$ if and only if $n(1-2^k\epsilon) < t \leq n\left(1-2^{k-1}\epsilon\right)$ with $k \in [L]$. Furthermore, define:

$$\begin{aligned}
\forall t \in [n], \forall i \in [m] : X_t^{*i} &= \left(A_{\sigma(t)} x^*_{\sigma(t)}\right)_i, \\
\forall t \in [n], \forall i \in [m] : X_t^{i} &= \left(A_{\sigma(t)} x_t\right)_i, \\
\forall t \in [n] : R_t^i &= \frac{\sum_{s=t}^n X_s^{*i}}{n-t+1} - \frac{1}{n}, \\
\forall t \in [n] : B_t &= \bigcap_{i=1}^m \left\{ \sigma \,\middle|\, R_t^i \leq \frac{1}{n} 2^{-\frac{\kappa(t)}{2}} \epsilon^{\frac{1}{2}} \right\}, \\
\forall k \in [L] : M_k^i &= \max_{n(1-2^k\epsilon) < t \leq n(1-2^{k-1}\epsilon)} R_t^i.
\end{aligned}$$

Note that $E\left[\frac{1}{n-t+1}\sum_{s=t}^n X_s^{*i}\right] \leq \frac{1}{n}$ for all $i \in [m]$ and $t \in [n]$. The event B_t consists of all the permutations for which $\max_i \frac{1}{n-t+1}\sum_{s=t}^n X_s^{*i}$ does not exceed $\frac{1}{n}$ by more than a small fraction of $\frac{1}{n}$. The purpose of the next two paragraphs is to bound the probability of $\cup_1^{n(1-\epsilon)} \bar{B}_t$. To do so, we first show that for all i, R_t^i is a martingale:

$$\begin{aligned}
E[R_t^i | R_{t-1}^i, \ldots, R_1^i] &= E\left[\frac{\sum_{s=t-1}^n X_s^{*i} - X^{*i}_{t-1}}{n-t+1} - \frac{1}{n} \,\middle|\, R_{t-1}^i, \ldots, R_1^i\right] \\
&= \frac{\sum_{s=t-1}^n X_s^{*i}}{n-t+1} - \frac{1}{n} - \frac{1}{n-t+1} E\left[X^{*i}_{t-1} \,\middle|\, R_{t-1}^i, \ldots, R_1^i\right] \\
&= \frac{\sum_{s=t-1}^n X_s^{*i}}{n-t+1} - \frac{1}{n} - \frac{\sum_{s=t-1}^n X_s^{*i}}{(n-t+1)(n-t+2)} = R_{t-1}^i.
\end{aligned}$$

By Doob's maximal inequality and Bernstein's inequality:

$$\forall k \in [L] : \quad \mathbb{P}\left(M_k^i > \frac{2^{-\frac{k}{2}}\epsilon^{\frac{1}{2}}}{n}\right) \leq \inf_{r \geq 0} E\left[\exp\left(r R^i_{n(1-2^{k-1}\epsilon)}\right)\right] \exp\left(-\frac{r 2^{-\frac{k}{2}}\epsilon}{n}\right) \leq \exp\left(\frac{-\epsilon^2}{6\gamma}\right), \tag{9}$$

$$\Rightarrow \quad \mathbb{P}\left(\bigcup_{t=1}^{n(1-\epsilon)} \bar{B}_t\right) \leq \sum_{k=1}^L \mathbb{P}\left(\bigcup_{n(1-2^k\epsilon)}^{n(1-2^{k-1}\epsilon)} \bar{B}_t\right) = \sum_{k=1}^L \mathbb{P}\left(\max_i M_k^i > \frac{2^{-\frac{k}{2}}\epsilon^{\frac{1}{2}}}{n}\right) \leq mL \exp\left(\frac{-\epsilon^2}{6\gamma}\right). \tag{10}$$

Theorem 1. *Suppose x_t, $1 \leq t \leq n$ is given by algorithm 1. Let $T = n(1-\epsilon)$. If $\epsilon \leq \frac{1}{4}$ and $\gamma \leq \frac{\epsilon^2}{12\log(m/\epsilon)}$, then $\mathbb{P}\left(\max_i \sum_{s=1}^{T} X_s^i > (1+2\epsilon)\right) \leq \epsilon$.*

Proof. Define:

$$\forall t \in [T] : \beta_t = \frac{\epsilon}{\gamma n}\left(1 + 2^{-\frac{\kappa(t)}{2}}\epsilon^{\frac{1}{2}}\right),$$
$$\forall t \in [T], \forall i \in [m] : \phi_i^t = \exp\left(\sum_{s=1}^{t} \nu\left(A_{\sigma(s)}x_s\right)_i - \beta_s\right),$$
$$\forall t \in [T] : \Phi^t = \sum_{i=1}^{m} \phi_i^t \prod_{s=1}^{t} 1_{B_s}.$$

For all $i \in [m]$, define $\phi_i^0 = 1$ and $\Phi^0 = \sum_{i=1}^{m} \phi_i^0 = m$. Let $\mathcal{A}_0 = \{\varnothing, \Pi\}$, and let $\mathcal{A}_t$ be the sigma algebra generated by $\sigma(1), \ldots, \sigma(t)$. We first show that $\{\Phi_t, \mathcal{A}_t\}$, $t \in \{0, 1, \ldots, T\}$ is a super-martingale:

$$\begin{aligned}
E\left[\Phi^t \middle| \mathcal{A}_{t-1}\right] &= E\left[\left(\sum_{i=1}^{m} \phi_i^{t-1} \exp\left(\nu X_t^i - \beta_t\right)\right)\prod_{s=1}^{t} 1_{B_s} \middle| \mathcal{A}_{t-1}\right] \\
&\leq E\left[\left(\sum_{i=1}^{m} \phi_i^{t-1}\left(1 + \frac{\epsilon}{\gamma}X_t^i\right)\exp\left(-\beta_t\right)\right)\prod_{s=1}^{t} 1_{B_s} \middle| \mathcal{A}_{t-1}\right] \\
&\leq E\left[\left(\sum_{i=1}^{m} \phi_i^{t-1}\left(1 + \frac{\epsilon}{\gamma}X_t^{*i}\right)\exp\left(-\beta_t\right)\right)\prod_{s=1}^{t} 1_{B_s} \middle| \mathcal{A}_{t-1}\right] \\
&= \sum_{i=1}^{m} \phi_i^{t-1} E\left[1 + \frac{\epsilon}{\gamma}X_t^{*i} \middle| \mathcal{A}_{t-1}\right]\exp\left(-\beta_t\right)\prod_{s=1}^{t} 1_{B_s} \\
&\leq \sum_{i=1}^{m} \phi_i^{t-1} 1_{\left\{R_t^i \leq \frac{1}{n}2^{-\kappa(t)/2}\epsilon^{1/2}\right\}}\exp\left(\frac{\epsilon}{\gamma}\left(\frac{1}{n} + R_t^i\right) - \beta_t\right)\prod_{s=1}^{t-1} 1_{B_s} \\
&\leq \sum_{i=1}^{m} \phi_i^{t-1}\prod_{s=1}^{t-1} 1_{B_s} = \Phi^{t-1}.
\end{aligned}$$

The first inequality follows from:

$$\forall \nu \in \mathbb{R}, \forall x \in [0,1] : \quad e^{\nu x} \leq 1 + x(e^{\nu} - 1). \tag{11}$$

The second inequality follows from the fact that

$$x_t \in \underset{z \in \Delta}{\operatorname{argmin}} \sum_{i=1}^{m} \exp\left(\nu \sum_{s=1}^{t-1}\left(A_{\sigma(s)}x_s\right)_i\right)\left(A_{\sigma(t)}z\right)_i,$$

and thus x_t as the minimizer achieves a smaller objective value than any other $x \in \Delta$ including $x^*_{\sigma(t)}$. Let $F = \cap_{s=1}^{T} B_s$. The previous result combined with Markov's inequality yields:

$$\begin{aligned}
\mathbb{P}\left(\left\{\max_i \sum_{s=1}^{T} X_s^i > \frac{T}{n}(1+3\epsilon)\right\} \cap F\right) &= \mathbb{P}\left(\left\{\max_i \phi_i^T > \exp\left(\nu\frac{T}{n}(1+3\epsilon) - \sum_{s=1}^{T}\beta_s\right)\right\} \cap F\right) \\
&\leq \mathbb{P}\left(\left\{\sum_{i=1}^{m} \phi_i^T > \exp\left(\nu\frac{T}{n}(1+3\epsilon) - \sum_{s=1}^{T}\beta_s\right)\right\} \cap F\right) \\
&\leq \mathbb{P}\left(\left\{\sum_{i=1}^{m} \phi_i^T \prod_{t=1}^{T} 1_{B_t} > \exp\left(\nu\frac{T}{n}(1+3\epsilon) - \sum_{s=1}^{T}\beta_s\right)\right\}\right) \\
&\leq E\left[\Phi^T\right] \exp\left(\frac{T}{n\gamma}\left(\epsilon(1+2\epsilon) - (1+3\epsilon)\nu\gamma\right)\right) \\
&\leq E\left[\Phi^0\right] \exp\left(\frac{T(1+2\epsilon)}{n\gamma}\left(\epsilon - \left(1+\frac{\epsilon}{1+2\epsilon}\right)\log(1+\epsilon)\right)\right) \\
&\leq m\exp\left(\frac{-(1-\epsilon)(1+2\epsilon)\epsilon^2}{6\gamma(1+\epsilon/3)}\right) \leq m\exp\left(\frac{-\epsilon^2}{6\gamma}\right). \qquad (12)
\end{aligned}$$

In the last line, we used the following inequality:

$$\forall u > 0, \forall \alpha \in \left[\frac{1}{2}, 1\right]: \quad u - (1+\alpha u)\log(1+u) \leq \frac{-(\alpha - 1/2)u^2}{1+u/3}. \qquad (13)$$

Using the bounds given by (10) and (12), we can conclude that:

$$\begin{aligned}
\mathbb{P}\left(\max_i \sum_{s=1}^{T} X_s^i > (1+2\epsilon)\right) &\leq \mathbb{P}\left(\max_i \sum_{s=1}^{T} X_s^i > \frac{T}{n}(1+3\epsilon)\right) \\
&\leq \mathbb{P}\left(\left\{\max_i \sum_{s=1}^{T} X_s^i > \frac{T}{n}(1+3\epsilon)\right\} \cap F\right) + \mathbb{P}\left(\bar{F}\right) \\
&\leq m(L+1)\exp\left(\frac{-\epsilon^2}{6\gamma}\right) \leq \epsilon,
\end{aligned}$$

where the second inequality follows from the fact that for any two events E and F, $\mathbb{P}(E) = \mathbb{P}(E \cap F) + \mathbb{P}(E \cap \bar{F}) \leq \mathbb{P}(E \cap F) + \mathbb{P}(\bar{F})$. In the last line, we used the fact that $L = \log_2(\frac{1}{\epsilon}) \leq \log_2(e)\epsilon^{-\frac{1}{e}}$ alongside the assumption that $\epsilon \leq \frac{1}{4}$. □

In the next section, we introduce ESA in its full generality and provide a competitive ratio analysis for ESA under the random permutation model.

4 General problem

Recall that the offline optimization problem is as follows:

$$P^* = \sup_{x_1,\ldots,x_n \in \mathbb{R}^k} \quad \sum_{t=1}^{n} f_t(x_t) + \psi\left(\sum_{t=1}^{n} A_t x_t\right), \qquad (\mathcal{P}1)$$

where for all $t \in [n]$, $A_t \in \mathbb{R}_+^{m\times k}$, $f_t \in \mathcal{G}$, and $\psi(u) = I_{u\leq b}$ for some $b \in \mathbb{R}_{++}^m$. Here $\mathcal{G}$ is a set of proper, concave, and upper semi-continuous functions on $\mathbb{R}^k$ with bounded super-level sets such that for all $f \in \mathcal{G}$, $\mathrm{dom} f \subset \mathbb{R}_+^k$ and $f(0) = 0$.

The Fenchel dual program is:

$$D^* = \inf_{y\in\mathbb{R}^m} \quad \sum_{t=1}^{n} -f_t^*(A_t^\top y) - \psi^*(-y). \qquad (\mathcal{D}1)$$

A pair of primal and dual variables $((x_1^*, x_2^*, \ldots, x_n^*), y^*)$ is an optimal pair when:

$$y^* \in -\partial\psi\left(\sum_{t=1}^{n} A_t x_t^*\right), \tag{14}$$

$$\forall t \in [n]: \quad A_t^\top y^* \in \partial f_t\left(x_t^*\right). \tag{15}$$

By our assumption on the functions involved such a pair exists and $0 \leq P^* = D^* < +\infty$. We assume that $P^* > 0$. Let γ be such that:

$$\forall t \in [n]: \quad \gamma \geq \max\left(\left\{ \frac{(A_t x)_i}{b_i} \middle| \, i \in [m], f_t(x) \geq 0 \right\} \cup \left\{ \frac{f_t(x)}{P^*} \middle| \, x \in \mathrm{dom} f_t \right\}\right). \tag{16}$$

Note that for all $t \in [n]$:

$$\begin{aligned} &0 \leq f_t(x_t^*), \\ &0 \leq -f_t^*(A_t^\top y^*) = f_t(x_t^*) - {y^*}^\top A_t x_t^* \leq \gamma P^*. \end{aligned}$$

ESA is designed to solve the following problem which is equivalent to $(\mathcal{P}1)$:

$$\underset{x_1,\ldots,x_n \in \mathbb{R}^k}{\text{maximize}} \quad \hat{\psi}\left(-\sum_{t=1}^{n} f_t(x_t)\right) + \psi\left(\sum_{t=1}^{n} A_t x_t\right), \tag{$\mathcal{P}2$}$$

where $\hat{\psi}(v) = I_{v \leq -P^*}$. Since P^* cannot be calculated without having access to $\left(\left(f_1, A_1\right), \left(f_2, A_2\right), \ldots, \left(f_n, A_n\right)\right)$, ESA retains an estimate of P^* and updates that estimate in exponential intervals. Choose σ uniformly at random from Π. Fix $\epsilon \in (0,1)$ and let $h \in \{0, \ldots, L-1\}$. To estimate the optimal value using $(f_{\sigma(t)}, A_{\sigma(t)})$, $1 \leq t \leq 2^h n\epsilon$, consider the following optimization problem:

$$\begin{aligned} P_h = &\sup_{x_t \in \mathbb{R}^k \;\; t \in [2^h n\epsilon]} \quad \frac{1}{2^h \epsilon(1-\theta_h)} \sum_{t=1}^{2^h n\epsilon} f_{\sigma(t)}(x_t) + \psi\left(\frac{1}{2^h \epsilon(1+\theta_h)} \sum_{t=1}^{2^h n\epsilon} A_{\sigma(t)} x_t\right) \\ \leq &\inf_{y \in \mathbb{R}^m} \quad \frac{1}{2^h \epsilon(1-\theta_h)} \sum_{t=1}^{2^h n\epsilon} -f_{\sigma(t)}^*\left(A_{\sigma(t)}^\top y\right) - \psi^*\left(-\frac{1+\theta_h}{1-\theta_h} y\right), \end{aligned} \tag{17}$$

where $\theta_h = 2^{-\frac{h+1}{2}} \epsilon^{\frac{1}{2}}$. In order to compare P_h with P^*, we define:

$$\begin{aligned} \tilde{P}_h &= \frac{1}{2^h \epsilon(1-\theta_h)} \sum_{t=1}^{2^h n\epsilon} f_{\sigma(t)}(x_{\sigma(t)}^*) + \psi\left(\frac{1}{2^h \epsilon(1+\theta_h)} \sum_{t=1}^{2^h n\epsilon} A_{\sigma(t)} x_{\sigma(t)}^*\right), \\ \tilde{D}_h &= \frac{1}{2^h \epsilon(1-\theta_h)} \sum_{t=1}^{2^h n\epsilon} -f_{\sigma(t)}^*\left(A_{\sigma(t)}^\top y^*\right) - \psi^*\left(-\frac{1+\theta_h}{1-\theta_h} y^*\right). \end{aligned}$$

Note that $\tilde{D}_h \geq P_h \geq \tilde{P}_h$. Now using the previous fact alongside Bernstein's inequality, we show that P_h is close to P^* with high probability.

$$\begin{aligned} \mathbb{P}\left(P_h < P^*\right) &\leq \mathbb{P}\left(\tilde{P}_h < P^*\right) \\ &\leq \mathbb{P}\left(\sum_{t=1}^{2^h n\epsilon} f_{\sigma(t)}(x_{\sigma(t)}^*) < 2^h \epsilon(1-\theta_h) P^*\right) + \sum_{i=1}^{m} \mathbb{P}\left(\sum_{t=1}^{2^h n\epsilon} (A_{\sigma(t)} x_{\sigma(t)}^*)_i > 2^h \epsilon(1+\theta_h) b_i\right) \\ &\leq \exp\left(\frac{-\epsilon^2}{4\gamma}\right) + m \exp\left(\frac{-\epsilon^2}{6\gamma}\right), \end{aligned} \tag{18}$$

$$\mathbb{P}\left(P_h > \frac{1+2\theta_h}{1-\theta_h}P^*\right) \leq \mathbb{P}\left(\tilde{D}_h > \frac{1+2\theta_h}{1-\theta_h}P^*\right)$$
$$\leq \mathbb{P}\left(\sum_{t=1}^{2^h n\epsilon} -f^*_{\sigma(t)}(A^\top_{\sigma(t)}y^*) > -2^h\epsilon(1+\theta_h)\sum_{t=1}^{n} f^*_{\sigma(t)}(A^\top_{\sigma(t)}y^*) + 2^h\epsilon\theta_h P^*\right) \leq \exp\left(\frac{-\epsilon^2}{6\gamma}\right). \quad (19)$$

Let κ be defined as in section 3 and let $\eta : \{n\epsilon+1, n\epsilon+2, \ldots, n\} \mapsto [L]$ be defined as $\eta(t) = \kappa(n-t+1)$. Equivalently, $\eta(t) = h$ if and only if $n2^{h-1}\epsilon < t \leq n2^h\epsilon$ with $h \in [L]$. For all $t \in \{n\epsilon+1, n\epsilon+2, \ldots, n\}$, let $q_t = P_{\eta(t)-1}$ if $P_{\eta(t)-1} \neq 0$. To avoid division by zero, we set $q_t = \epsilon$ if $P_{\eta(t)-1} = 0$. Note that since $P^* > 0$, the probability of $P_{\eta(t)-1} = 0$ is bounded by (18). Now we define parameters that will be used in ESA. For all $t \in \{n\epsilon+1, n\epsilon+2, \ldots, n(1-\epsilon)\}$, define:

$$\alpha_t = \frac{1-2^{-\frac{\eta(t)}{2}}\epsilon^{\frac{1}{2}}}{1+2^{-\frac{\eta(t)}{2}+1}\epsilon^{\frac{1}{2}}},\ \beta_t = \frac{\epsilon}{\gamma n}\left(1+2^{-\frac{\kappa(t)}{2}}\epsilon^{\frac{1}{2}}\right),\ \beta'_t = \frac{\epsilon}{\gamma n}\left(1-2^{-\frac{\kappa(t)}{2}}\epsilon^{\frac{1}{2}}\right)\alpha_t,$$

while for all $t \in \{n(1-\epsilon)+1, \ldots, n+1\}$, let $\beta_t = \beta_{n(1-\epsilon)}$ and $\beta'_t = \beta'_{n(1-\epsilon)}$.

Algorithm 2 Exponentiated Subgradient Algorithm (ESA)

Require: $\epsilon \in (0,1)$, $\gamma > 0$
  $\nu = \frac{1}{\gamma}\log(1+\epsilon)$, $\nu' = \frac{-1}{\gamma}\log(1-\epsilon)$
  for $t = 1 : n\epsilon$ **do**
    $\check{x}_t = 0$
  end for
  $y'_{n\epsilon+1} = \exp(\beta'_{n\epsilon+1})$, $y_{n\epsilon+1,i} = \exp(-\beta_{n\epsilon+1})/m \quad (\forall i \in [m])$
  for $t = n\epsilon+1 : n$ **do**
    $x_t \in \operatorname{argmin}_z \sum_{i=1}^{m} \frac{y_{t,i}}{b_i}\left(A_{\sigma(t)}z\right)_i - \frac{y'_t}{q_t}f_{\sigma(t)}(z)$
    $y_{t+1,i} = y_{t,i}\exp\left(\frac{\nu}{b_i}\left(A_{\sigma(t)}x_t\right)_i - \beta_{t+1}\right) \quad (\forall i \in [m])$
    $y'_{t+1} = y'_t\exp\left(\frac{-\nu'}{q_t}f_{\sigma(t)}(x_t) + \beta'_{t+1}\right)$
    if $\sum_{s=n\epsilon+1}^{t} A_{\sigma(s)}x_s \leq b$ **then**
      $\check{x}_t = x_t$
    else
      $\check{x}_t = \mathbf{0}$
    end if
  end for

Algorithm (2) describes ESA applied to ($\mathcal{P}1$). In section 4.1, we represent each update of (y_t, y'_t) as an exponentiated subgradient step. Let $x_t, 1 \leq t \leq n$ be given by algorithm (2). For all $t \in [n]$, define:

$$y''_t = \frac{q_t}{y'_t}\left[y_{t,1}/b_1,\ y_{t,2}/b_2,\ \ldots,\ y_{t,m}/b_m\right]^\top.$$

Note that $x_t \in \operatorname{argmin} {y''_t}^\top A_{\sigma(t)}z - f_{\sigma(t)}(z)$ if and only if $A^\top_{\sigma(t)}y''_t \in \partial f_{\sigma(t)}(x_t)$ if and only if $x_t \in \partial f^*_{\sigma(t)}\left(A^\top_{\sigma(t)}y''_t\right)$. For special cases of $f_{\sigma(t)}$, this assignment rule can have a simple form; for example, if $k = 1$ and $f_{\sigma(t)}(x) = \tilde{f}(x) + I_{x\in[0,1]}$ with $\tilde{f}$ differentiable, then x_t can be chosen as:

$$x_t = \begin{cases} 0 & \tilde{f}'(0) \leq A^\top_{\sigma(t)}y''_t, \\ 1 & \tilde{f}'(1) > A^\top_{\sigma(t)}y''_t, \\ l(A^\top_{\sigma(t)}y''_t) & \text{otherwise,} \end{cases}$$

where $l(v) = \min\{x | v \geq \tilde{f}'(x)\}$. Similar to the analysis in section 3, we need to derive concentration inequalities for the running sums of form $\sum_{s=t}^{n}\left(A_{\sigma(s)}x^*_{\sigma(s)}\right)_i$ and $\sum_{s=t}^{n} f_{\sigma(s)}\left(x^*_{\sigma(s)}\right)$. Define:

$$
\begin{aligned}
\forall i \in [m],\, \forall t \in [n]: \quad & X_t^{*i} = \left(A_{\sigma(t)} x^*_{\sigma(t)}\right)_i, \\
\forall i \in [m],\, \forall t \in [n]: \quad & R_t^i = \frac{\sum_{s=t}^n X_s^{*i}}{n-t+1} - \frac{b_i}{n}, \\
\forall t \in [n]: \quad & S_t = \frac{\sum_{s=t}^n f_{\sigma(s)}\left(x^*_{\sigma(s)}\right)}{n-t+1} - \frac{P^*}{n}, \\
\forall t \in [n]: \quad & C_t = \left\{ \sigma \,\middle|\, S_t \geq -\frac{1}{n} P^* 2^{-\frac{\kappa(t)}{2}} \epsilon^{\frac{1}{2}} \right\}, \\
\forall t \in [n]: \quad & B_t = \bigcap_{i=1}^m \left\{ \sigma \,\middle|\, R_t^i \leq \frac{1}{n} b_i 2^{-\frac{\kappa(t)}{2}} \epsilon^{\frac{1}{2}} \right\}, \\
\forall i \in [m],\, \forall k \in [L]: \quad & M_k^i = \max_{n(1-2^k\epsilon) < t \leq n(1-2^{k-1}\epsilon)} R_t^i, \\
\forall k \in [L]: \quad & N_k = \max_{n(1-2^k\epsilon) < t \leq n(1-2^{k-1}\epsilon)} -S_t.
\end{aligned}
$$

As in section 3, for all i, R_t^i is a martingale. Similarly, it can be shown that S_t is a martingale. By Doob's maximal inequality and Bernstein's inequality:

$$
\begin{aligned}
\forall i \in [m],\, \forall k \in [L]: \quad & \mathbb{P}\left(M_k^i > \frac{1}{n} b_i 2^{-\frac{k}{2}} \epsilon^{\frac{1}{2}}\right) \leq \inf_{r \geq 0} E\left[\exp\left(r R^i_{n(1-2^{k-1}\epsilon)}\right)\right] \exp\left(-\frac{1}{n} r b_i 2^{-\frac{k}{2}} \epsilon^{\frac{1}{2}}\right) \leq \exp\left(\frac{-\epsilon^2}{6\gamma}\right) \\
\Rightarrow \quad & \mathbb{P}\left(\bigcup_{t=1}^{n(1-\epsilon)} \bar{B}_t\right) \leq \sum_{k=1}^L \mathbb{P}\left(\bigcup_{n(1-2^k\epsilon)}^{n(1-2^{k-1}\epsilon)} \bar{B}_t\right) = \sum_{k=1}^L \mathbb{P}\left(\max_i \frac{M_k^i}{b_i} > \frac{1}{n} 2^{-\frac{k}{2}} \epsilon^{\frac{1}{2}}\right) \leq mL \exp\left(\frac{-\epsilon^2}{6\gamma}\right). \quad (20)
\end{aligned}
$$

$$
\begin{aligned}
\forall k \in [L]: \quad & \mathbb{P}\left(N_k > \frac{1}{n} P^* 2^{-\frac{k}{2}} \epsilon^{\frac{1}{2}}\right) \leq \inf_{r \geq 0} E\left[\exp\left(-r S_{n(1-2^{k-1}\epsilon)}\right)\right] \exp\left(-\frac{1}{n} r P^* 2^{-\frac{k}{2}} \epsilon^{\frac{1}{2}}\right) \leq \exp\left(\frac{-\epsilon^2}{4\gamma}\right) \\
\Rightarrow \quad & \mathbb{P}\left(\bigcup_{t=1}^{n(1-\epsilon)} \bar{C}_t\right) \leq \sum_{k=1}^L \mathbb{P}\left(\bigcup_{n(1-2^k\epsilon)}^{n(1-2^{k-1}\epsilon)} \bar{C}_t\right) = \sum_{k=1}^L \mathbb{P}\left(N_k > \frac{1}{n} P^* 2^{-\frac{k}{2}} \epsilon^{\frac{1}{2}}\right) \leq L \exp\left(\frac{-\epsilon^2}{4\gamma}\right). \quad (21)
\end{aligned}
$$

Theorem 2. *Suppose $\check{x}_t$, $1 \leq t \leq n$ is given by algorithm (2). If $\epsilon \leq \frac{1}{12}$ and $\gamma \leq \frac{\epsilon^2}{13 \log(m/\epsilon)}$, then:*

$$
E\left[\sum_{t=1}^n f_t(\check{x}_t) + \psi\left(\sum_{t=1}^n A_t \check{x}_t\right)\right] \geq (1 - 12\epsilon) P^*. \quad (22)
$$

Remark. *The conclusion of the theorem can be restated as:*

$$
C_{\mathcal{D}_2} \geq 1 - 12\epsilon, \quad (23)
$$

where $\mathcal{H}$ is the set of all $((f_1, A_1), (f_2, A_2), \ldots, (f_n, A_n)) \in \left(\mathcal{G} \times \mathbb{R}_+^{m \times k}\right)^n$ that satisfy (16) with $\gamma = \frac{\epsilon^2}{13 \log(m/\epsilon)}$, and with $P^* > 0$.

Proof. Here we state the main steps of the proof. The details are presented in appendix A. Let $T = n(1 - 2\epsilon)$. For all $t \in \{n\epsilon + 1, n\epsilon + 2, \ldots, T\}$, define:

$$
\begin{aligned}
& F_t = C_t \cap B_t \cap \{P^* \leq q_t \leq P^*/\alpha_t\}, \\
\forall i \in [m]: \quad & \phi_i^t = \exp\left(\sum_{s=n\epsilon+1}^{t} \frac{\nu}{b_i}\left(A_{\sigma(s)}x_s\right)_i - \beta_s\right), \\
& \chi^t = \exp\left(\sum_{s=n\epsilon+1}^{t} -\frac{\nu'}{q_s} f_{\sigma(s)}(x_s) + \beta'_s\right), \\
& \Phi^t = \left(\sum_{i=1}^{m} \phi_i^t + m\chi^t\right) \prod_{s=n\epsilon+1}^{t} 1_{F_s}.
\end{aligned}
$$

For all $i \in [m]$, define $\phi_i^{n\epsilon} = 1$, $\chi^{n\epsilon} = 1$, and $\Phi^{n\epsilon} = \sum_{i=1}^{m} \phi_i^{n\epsilon} + m\chi^{n\epsilon} = 2m$. Let $\mathcal{A}_{n\epsilon} = \{\varnothing, \Pi\}$, and for all $t \in \{n\epsilon+1, n\epsilon+2, \ldots, T\}$, let $\mathcal{A}_t$ be the sigma algebra generated by $\sigma(1), \ldots, \sigma(t)$. We show in appendix A that $\{\Phi_t, \mathcal{A}_t\}$, $t \in \{n\epsilon, n\epsilon+1, \ldots, T\}$, is a super-martingale, which yields:

$$
E\left[\Phi^T\right] \leq E\left[\Phi^{n\epsilon}\right] = 2m. \tag{24}
$$

Let $B = \left\{\frac{1}{P^*}\sum_{s=n\epsilon+1}^{T} f_{\sigma(s)}(x_s) < (1-11\epsilon)\right\}$, $C = \left\{\max_i \frac{1}{b_i}\sum_{s=n\epsilon+1}^{T}\left(A_{\sigma(s)}x_s\right)_i > 1\right\}$, and $F = \cap_{s=1}^{T} F_s$. Using Markov's inequality, we can derive:

$$
\begin{aligned}
\mathbb{P}\left(B \cap F\right) &\leq E\left[\chi^T 1_F\right] \exp\left(-\frac{\epsilon^2}{7\gamma}\right), \\
\mathbb{P}\left(C \cap F\right) &\leq E\left[\sum_{i=1}^{m} \phi_i^T 1_F\right] \exp\left(-\frac{\epsilon^2}{4\gamma}\right).
\end{aligned}
$$

Combining the previous bounds with (18), (19), (20) and (21), we get:

$$
\begin{aligned}
\mathbb{P}\left(B \cup C\right) \leq \mathbb{P}\left(B \cap F\right) + \mathbb{P}\left(C \cap F\right) + \mathbb{P}\left(\bar{F}\right) &\leq E\left[\chi^T 1_F\right] \exp\left(-\frac{\epsilon^2}{7\gamma}\right) + E\left[\sum_{i=1}^{m} \phi_i^T 1_F\right] \exp\left(-\frac{\epsilon^2}{4\gamma}\right) \\
&\quad + 2L\exp\left(\frac{-\epsilon^2}{4\gamma}\right) + (2m+1)L\exp\left(\frac{-\epsilon^2}{6\gamma}\right) \\
&\leq E\left[\Phi^T\right] \exp\left(-\frac{\epsilon^2}{7\gamma}\right) + 2L\exp\left(\frac{-\epsilon^2}{4\gamma}\right) + (2m+1)L\exp\left(\frac{-\epsilon^2}{6\gamma}\right) \\
&\leq 2m\exp\left(-\frac{\epsilon^2}{7\gamma}\right) + 2L\exp\left(\frac{-\epsilon^2}{4\gamma}\right) + (2m+1)L\exp\left(\frac{-\epsilon^2}{6\gamma}\right) \leq \epsilon,
\end{aligned}
$$

where in the last line we used the fact that $L = \log_2(\frac{1}{\epsilon}) \leq \log_2(e)\epsilon^{-\frac{1}{e}}$ alongside the assumption that $\epsilon \leq \frac{1}{12}$. As a result, we get:

$$
\begin{aligned}
E\left[\sum_{t=1}^{n} f_t(\check{x}_t) + \psi\left(\sum_{t=1}^{n} A_t\check{x}_t\right)\right] &\geq E\left[\sum_{t=n\epsilon+1}^{T} f_t(x_t) + \psi\left(\sum_{t=n\epsilon+1}^{T} A_t x_t\right) \middle| \bar{B} \cap \bar{C}\right] \mathbb{P}\left(\bar{B} \cap \bar{C}\right) \\
&\geq (1-11\epsilon)(1-\epsilon)P^* \geq (1-12\epsilon)P^*.
\end{aligned}
$$

□

4.1 Connection with subgradient methods

In this section, we discuss how the updates of (y_t, y'_t) in algorithm (2) are equivalent to online exponentiated subgradient steps and point to the similarities and differences between ESA, DLA and KRTV. Recall that a pair $\left((x_1^*, x_2^*, \ldots, x_n^*), y^*\right)$ is an optimal primal-dual pair for $(\mathcal{P}1)$ if and only if:

$$y^* \in -\partial\psi\left(\sum_{t=1}^{n} A_t x_t^*\right), \tag{25}$$

$$\forall t \in [n]: \quad A_t^\top y^* \in \partial f_t\left(x_t^*\right). \tag{26}$$

Note that (26) is decoupled in t. Therefore, to assign x_t, an online algorithm can find y_t, an estimate of y^*, and then choose x_t such that:

$$x_t \in \partial f^*_{\sigma(t)}\left(A^\top_{\sigma(t)} y_t\right) \leftrightarrow A^\top_{\sigma(t)} y_t \in \partial f_{\sigma(t)}\left(x_t\right). \tag{27}$$

ESA and DLA are two examples of online algorithms that choose x_t according to (27). KRTV can also be transformed into an algorithm that chooses x_t according to (27) (Both DLA and KRTV are proposed for online LP. Here we consider extensions of those algorithms to $(\mathcal{P}1)$). However, these algorithms update y_t differently. When $t = n\epsilon 2^h$ for some $h \in \{0, 1, \ldots, L-1\}$, DLA sets y_t to be:

$$y_t \in \underset{y\in\mathbb{R}^m}{\operatorname{argmin}} \quad \frac{1}{2^h\epsilon(1-2^{-\frac{h}{2}}\epsilon^{\frac{1}{2}})} \sum_{s=1}^{t} -f^*_{\sigma(s)}\left(A^\top_{\sigma(s)} y\right) - \psi^*\left(-y\right), \tag{28}$$

then for all $t \in \{n\epsilon 2^h + 1, n\epsilon 2^h + 2, \ldots, n\epsilon 2^{h+1}\}$ sets $y_t = y_{n\epsilon 2^h}$. Although the dual variables do not explicitly appear in the description of KRTV in [KRTV14], this algorithm can be represented as an extreme case that updates y_t by solving (28) (without $(1-2^{-\frac{h}{2}}\epsilon^{\frac{1}{2}})$) for every $t \in [n]$. ESA also updates y_t for every $t \in [n]$. However, the updates are very simple and computationally very cheap. Consider the following problem which is dual to $(\mathcal{P}2)$:

$$\underset{y\in\mathbb{R}^m,\, y'\in\mathbb{R}}{\text{minimize}} \quad -\psi^*\left(-y\right) - \hat{\psi}^*\left(-y'\right) + \sum_{t=1}^{n} \sup_{x\in\mathbb{R}^k} y' f_t(x) - y^\top A_t x. \tag{$\mathcal{D}2$}$$

Note that $\psi^*\left(u\right) = b^T u$ and $\hat{\psi}^*\left(v\right) = -P^* v$ for all $u \in \mathbb{R}^m_+$ and $v \in \mathbb{R}_+$. Using the change of variables $y_i \to y_i/b_i$ and $y \to y/P^*$, we can rewrite $(\mathcal{D}2)$ as:

$$\underset{y\in\mathbb{R}^m_+,\, y'\in\mathbb{R}_+}{\text{minimize}} \quad \mathbf{1}^\top y - y' + \sum_{t=1}^{n} H_t(y, y'), \tag{29}$$

where $H_t(y,y') = \sup_{x\in\mathbb{R}^k} \frac{y'}{P^*} f_t(x) - \sum_{i=1}^m \frac{y_i}{b_i}\left(A_t x\right)_i$. By replacing P^* with q_t, we can define:

$$\tilde{H}_{\sigma(t)}(y,y') = \sup_{x\in\mathbb{R}^k} \frac{y'}{q_t} f_{\sigma(t)}(x) - \sum_{i=1}^{m} \frac{y_i}{b_i}\left(A_{\sigma(t)} x\right)_i.$$

Let x_t, $n\epsilon < t \le n$ be given by ESA. Since for all $t \in \{n\epsilon+1, n\epsilon+2, \ldots, n\}$, $x_t \in \partial f^*_{\sigma(t)}(A^\top_{\sigma(t)} y''_t)$, we get:

$$\left[\ \tfrac{-1}{b_1}\left(A_{\sigma(t)}x_t\right)_1,\ \ \tfrac{-1}{b_2}\left(A_{\sigma(t)}x_t\right)_2,\ \ \ldots\ \ \tfrac{-1}{b_m}\left(A_{\sigma(t)}x_t\right)_m,\ \ \tfrac{1}{q_t} f_{\sigma(t)}(x_t)\ \right]^\top \in \partial\tilde{H}_{\sigma(t)}(y_t, y'_t). \tag{30}$$

For all $t \in \{n\epsilon+1, \ldots, n-1\}$, ESA updates (y_t, y'_t) via an online exponentiated step toward the negative direction of the subgradient of $\tilde{H}_{\sigma(t)}(y_t, y'_t)$ given in (30):

$$\forall i \in [m]: \quad (y_{t+1})_i = y_t \exp\left(\frac{\nu}{b_i}\left(A_{\sigma(t)}x_t\right)_i - \beta_{t+1}\right),$$
$$y'_{t+1} = y'_t \exp\left(-\frac{\nu'}{q_t} f_{\sigma(t)}(x_t) + \beta'_{t+1}\right),$$

where ν and ν' are the step sizes while $\exp\left(-\beta_{t+1}\right)$ and $\exp\left(\beta'_{t+1}\right)$ are the normalization terms.

Table 1: Comparison between OLA, DLA, ESA, KRTV, KRTV5 when $c \in \{1200, 300, 60, 30\}$

		$n = 3630, c = 1200$		$n = 915, c = 300$		$n = 189, c = 60$		$n = 96, c = 30$	
Algorithm	ϵ	CR	time (s)	CR	time (s)	CR	time (s)	CR	time (s)
OLA	0.1	0.88	0.060	0.79	0.020	0.72	0.012	0.65	0.011
	0.05	0.91	0.064	0.83	0.020	0.69	0.022	0.67	0.011
DLA	0.1	0.89	0.112	0.74	0.05	0.79	0.036	0.76	0.033
	0.05	0.94	0.148	0.92	0.065	0.84	0.081	0.80	0.043
ESA	0.1	0.86	0.173	0.87	0.068	0.90	0.041	0.86	0.036
	0.05	0.93	0.205	0.95	0.084	0.92	0.093	0.87	0.046
KRTV	n/a	0.97	132.636	0.93	17.485	0.88	4.272	0.84	1.125
KRTV5	n/a	0.98	27.828	0.94	3.493	0.90	0.831	0.87	0.221

5 Numerical experiment

In this section, we examine the performance of ESA and compare it with that of four other online LP algorithms proposed in the literature. The algorithms are *one-time learning algorithm (OLA)*, DLA, KRTV, and KRTV5. OLA is a simpler version of DLA introduced in [AWY09]. OLA computes the dual variable at $t = n\epsilon$ and uses the same dual variable to choose x_t for all $t \in \{n\epsilon + 1, n\epsilon + 2, \ldots, n\}$. As it was discussed in section 4.1, KRTV can be viewed as an algorithm that chooses x_t according to (27) and updates the dual variable at each step. This point of view motivated us to consider a variant of KRTV that updates the dual variable after every fixed number of steps. In KRTV5, the dual variable is updated every 5 steps.

Consider an online LP problem and let $M = \{\{(c_t, A_t) \,|1 \leq t \leq n\}\}$, where the double bracket notation denotes a multiset. we generate M based on the construction proposed in [AWY09] for deriving the necessary condition on γ. In this scheme of construction, $m = 2^d$ for some integer d. $b = c\mathbf{1}$ for some $c \in \mathbb{R}_{++}$. For all $l \in [m]$, let $l - 1 = (a_{l,1}, a_{l,2}, \ldots, a_{l,d})_2$. For all $i \in [d]$, define $v_i = [a_{1,i}, a_{2,i}, \ldots, a_{m,i}]^T$ and $w_i = \mathbf{1} - v_i$. Let $j_1, j_2, \ldots, j_m$ be m independent samples from $Binomial(\lceil \frac{2c}{d} \rceil, \frac{1}{2})$. M consists of:

1. $(4, v_i)$ with multiplicity $\lceil \frac{c}{d} \rceil$ for all $i \in [d]$
2. $(3, w_i)$ with multiplicity j_i for all $i \in [d]$
3. $(2, w_i)$ with multiplicity $\lceil \frac{1}{2}\sqrt{\frac{c}{d}} \rceil$ for all $i \in [d]$
4. $(1, w_i)$ with multiplicity $\lceil \frac{2c}{d} \rceil - j_i$ for all $i \in [d]$

Table 1 compares the competitive ratio and the runtime of all the algorithms when $d = 3$ and $c \in \{1200, 300, 60, 30\}$. The competitive ratio of the algorithms is estimated by averaging the performance for 500 permutations. The competitive ratios (CR) reported in Table 1 are the average of estimated competitive ratios for 10 independently generated instances of M. The runtime reports the average time it takes an algorithm to provide a solution for one permutation of a single instance of M. For this numerical experiment, we used Matlab R2014a and the linear programming solver of Matlab optimization toolbox. The runtimes are collected from a Linux machine with an Intel Core i7.4770 3.40GHz CPU and 8GB of RAM. In ESA, γ is set to be equal to $\frac{1}{c}$. Note that ϵ is as an input parameter for ESA, OLA, and DLA. For each $\epsilon \in \{0.1, 0.05\}$, we have considered two problems with $c = \frac{d}{\epsilon^2}$ and $c = \frac{d}{\epsilon}$. When $c = \frac{d}{\epsilon^2}$, the bid-to-budget ratio matches the necessary condition for achieving a competitive ratio of $1 - O(\epsilon)$. As the data presented in Table 1 suggest, ESA, OLA, and DLA are multiple orders of magnitude faster than KRTV and KRTV5. When $c \in \{300, 60, 30\}$, ESA with $\epsilon = 0.05$ achieves a competitive ratio that is higher than or equal to the rest of the algorithms. Note that ESA, OLA, and DLA choose $x_t = 0$ for $t \in [n\epsilon]$. Despite this fact, when $\epsilon = 0.05$ and $c = \frac{d}{0.05^2} = 1200$, the difference between the competitive ratio of ESA and KRTV is less than ϵ.

Appendix A: Details for the proof of Theorem 2

First, we show that $\{\Phi_t, \mathcal{A}_t\}$, $t \in \{n\epsilon, n\epsilon + 1, \ldots, T\}$ is a super-martingale.

$$
\begin{aligned}
E\left[\Phi^t\middle|\mathcal{A}_{t-1}\right] &= E\left[\left(\sum_{i=1}^m \phi_i^t + m\chi^t\right)\prod_{s=n\epsilon+1}^t 1_{F_s}\middle|\mathcal{A}_{t-1}\right] \\
&= E\left[\left(\sum_{i=1}^m \phi_i^{t-1}\exp\left(\frac{\nu}{b_i}\left(A_{\sigma(t)}x_t\right)_i - \beta_t\right) + m\chi^{t-1}\exp\left(-\frac{\nu'}{q_t}f_{\sigma(t)}\left(x_t\right)+\beta_t'\right)\right)\prod_{s=n\epsilon+1}^t 1_{F_s}\middle|\mathcal{A}_{t-1}\right] \\
&\leq E\left[\left(\sum_{i=1}^m \phi_i^{t-1}\left(1+\frac{\epsilon}{\gamma b_i}\left(A_{\sigma(t)}x_t\right)_i\right)\exp\left(-\beta_t\right) + m\chi^{t-1}\left(1-\frac{\epsilon}{\gamma q_t}f_{\sigma(t)}\left(x_t\right)\right)\exp\left(\beta_t'\right)\right)\prod_{s=n\epsilon+1}^t 1_{F_s}\middle|\mathcal{A}_{t-1}\right] \\
&\leq E\left[\left(\sum_{i=1}^m \phi_i^{t-1}\left(1+\frac{\epsilon}{\gamma b_i}\left(A_{\sigma(t)}x^*_{\sigma(t)}\right)_i\right)\exp\left(-\beta_t\right) + m\chi^{t-1}\left(1-\frac{\epsilon}{\gamma q_t}f_{\sigma(t)}\left(x^*_{\sigma(t)}\right)\right)\exp\left(\beta_t'\right)\right)\prod_{s=n\epsilon+1}^t 1_{F_s}\middle|\mathcal{A}_{t-1}\right] \\
&= \sum_{i=1}^m \phi_i^{t-1}E\left[\left(1+\frac{\epsilon\left(A_{\sigma(t)}x^*_{\sigma(t)}\right)_i}{\gamma b_i}\right)\exp\left(-\beta_t\right)\middle|\mathcal{A}_{t-1}\right]\prod_{s=n\epsilon+1}^t 1_{F_s} \\
&\quad + m\chi^{t-1}E\left[\left(1-\frac{\epsilon f_{\sigma(t)}\left(x^*_{\sigma(t)}\right)}{\gamma q_t}\right)\exp\left(\beta_t'\right)\middle|\mathcal{A}_{t-1}\right]\prod_{s=n\epsilon+1}^t 1_{F_s} \\
&= \sum_{i=1}^m \phi_i^{t-1}\left(1+\frac{\epsilon}{\gamma b_i}\left(R_t^i + b_i/n\right)\right)\exp\left(-\beta_t\right)\prod_{s=n\epsilon+1}^t 1_{F_s} + m\chi^{t-1}\left(1-\frac{\epsilon}{\gamma q_t}(S_t+P^*/n)\right)\exp\left(\beta_t'\right)\prod_{s=n\epsilon+1}^t 1_{F_s} \\
&\leq \left(\phi^{t-1}\exp\left(\beta_t-\beta_t\right) + m\chi^{t-1}\exp\left(-\beta_t'+\beta_t'\right)\right)\prod_{s=n\epsilon+1}^{t-1} 1_{F_s} = \Phi^{t-1}.
\end{aligned}
$$

The first inequality follows from (11); the second inequality follows from the fact that

$$
x_t \in \operatorname*{argmin}_z {y_t''}^\top A_{\sigma(t)}z - f_{\sigma(t)}\left(z\right);
$$

the last inequality follows from the definition of F_t. Using the definition of ϕ_i^t and χ^t alongside Markov's inequality we can derive:

$$
\begin{aligned}
\mathbb{P}\left(\left\{\frac{1}{P^*}\sum_{s=1}^T f_{\sigma(s)}(x_s) < (1-3\epsilon)\left(1-8\epsilon\right)\right\}\cap F\right) &\leq \mathbb{P}\left(\left\{\sum_{s=n\epsilon+1}^T \frac{f_{\sigma(s)}(x_s)}{q_s} < (1-3\epsilon)\left(1-8\epsilon\right)\right\}\cap F\right) \\
&= \mathbb{P}\left(\left\{\chi^T > \exp\left(\sum_{s=n\epsilon+1}^T \beta_s' - \nu'\left(1-3\epsilon\right)\left(1-8\epsilon\right)\right)\right\}\cap F\right) \\
&\leq \mathbb{P}\left(\chi^T 1_F > \exp\left(\sum_{s=n\epsilon+1}^T \beta_s' - \nu'\left(1-3\epsilon\right)\left(1-8\epsilon\right)\right)\right) \\
&\leq E\left[\chi^T 1_F\right]\exp\left(\frac{1-3\epsilon}{\gamma}\left(\gamma\nu'(1-8\epsilon) - \epsilon(1-7\epsilon)\right)\right) \\
&\leq E\left[\chi^T 1_F\right]\exp\left(\frac{(1-3\epsilon)(1-7\epsilon)}{\gamma}\left(\left(\frac{\epsilon}{1-7\epsilon}-1\right)\log\left(1-\epsilon\right)-\epsilon\right)\right) \\
&\leq E\left[\chi^T 1_F\right]\exp\left(-\frac{(1-3\epsilon)(1-7\epsilon)\epsilon^2}{2\gamma}\right) \leq E\left[\chi^T 1_F\right]\exp\left(-\frac{\epsilon^2}{7\gamma}\right),
\end{aligned}
\tag{31}
$$

where the last line follows from:

$$
\forall u\in[0,1]: \quad u + (1-u)\log(1-u) \geq \frac{1}{2}u^2. \tag{32}
$$

$$\mathbb{P}\left(\left\{\max_i \frac{1}{b_i}\sum_{s=1}^{T}\left(A_{\sigma(s)}x_s\right)_i > (1-3\epsilon)(1+3\epsilon)\right\}\cap F\right) = \mathbb{P}\left(\left\{\max_i \phi_i^T > \exp\left(\nu(1-3\epsilon)(1+3\epsilon) - \sum_{s=n\epsilon+1}^{T}\beta_s\right)\right\}\cap F\right)$$
$$\leq \mathbb{P}\left(\max_i \phi_i^T 1_F > \exp\left(\nu(1-3\epsilon)(1+3\epsilon) - \sum_{s=n\epsilon+1}^{T}\beta_s\right)\right)$$
$$\leq \mathbb{P}\left(\sum_{i=1}^{m}\phi_i^T 1_F > \exp\left(\nu(1-3\epsilon)(1+3\epsilon) - \sum_{s=n\epsilon+1}^{T}\beta_s\right)\right)$$
$$\leq E\left[\sum_{i=1}^{m}\phi_i^T 1_F\right]\exp\left(\frac{1-3\epsilon}{\gamma}\left(\epsilon(1+2\epsilon) - \gamma\nu(1+3\epsilon)\right)\right)$$
$$\leq E\left[\sum_{i=1}^{m}\phi_i^T 1_F\right]\exp\left(\frac{(1-3\epsilon)(1+2\epsilon)}{\gamma}\left(\epsilon - \left(1+\frac{\epsilon}{1+2\epsilon}\right)\log(1+\epsilon)\right)\right)$$
$$\leq E\left[\sum_{i=1}^{m}\phi_i^T 1_F\right]\exp\left(-\frac{(1-3\epsilon)(1+2\epsilon)\epsilon^2}{3\gamma(1+\epsilon/3)}\right) \leq E\left[\sum_{i=1}^{m}\phi_i^T 1_F\right]\exp\left(-\frac{\epsilon^2}{4\gamma}\right), \tag{33}$$

where the last line follows from (13).

References

[AD14] Shipra Agrawal and Nikhil R Devanur. Fast algorithms for online stochastic convex programming. *arXiv preprint arXiv:1410.7596*, 2014.

[AWY09] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. A dynamic near-optimal algorithm for online linear programming. *arXiv preprint arXiv:0911.2974*, 2009.

[BIKK07] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. A knapsack secretary problem with applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 16–28. Springer, 2007.

[BIKK08] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exch.*, 7(2):7:1–7:11, June 2008.

[BJN07] Niv Buchbinder, Kamal Jain, and Joseph Seffi Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *Algorithms–ESA 2007*, pages 253–264. Springer, 2007.

[BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford University Press, 2013.

[DH09] Nikhil R. Devanur and Thomas P. Hayes. The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, EC '09, pages 71–78, New York, NY, USA, 2009. ACM.

[DJ12] Nikhil R Devanur and Kamal Jain. Online matching with concave returns. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 137–144. ACM, 2012.

[DJSW11] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 29–38. ACM, 2011.

[FHK+10] Jon Feldman, Monika Henzinger, Nitish Korula, Vahab S. Mirrokni, and Cliff Stein. Online stochastic packing applied to display ad allocation. In *Proceedings of the 18th Annual European Conference on Algorithms: Part I*, ESA'10, pages 182–194, Berlin, Heidelberg, 2010. Springer-Verlag.

[FMMM09] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and S Muthukrishnan. Online stochastic matching: Beating 1-1/e. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 117–126. IEEE, 2009.

[GM14] Anupam Gupta and Marco Molinaro. How the experts algorithm can help solve lps online. *arXiv preprint arXiv:1407.5298*, 2014.

[JL12] Patrick Jaillet and Xin Lu. Near-optimal online algorithms for dynamic resource allocation problems. *arXiv preprint arXiv:1208.2596*, 2012.

[JL13] Patrick Jaillet and Xin Lu. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research*, 2013.

[Kle05] Robert Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 630–631. Society for Industrial and Applied Mathematics, 2005.

[KMT11] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 587–596. ACM, 2011.

[KP09] Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *Automata, Languages and Programming*, pages 508–520. Springer, 2009.

[KRTV13] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Algorithms–ESA 2013*, pages 589–600. Springer, 2013.

[KRTV14] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 303–312, New York, NY, USA, 2014. ACM.

[KVV90] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 352–358. ACM, 1990.

[Meh13] Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013.

[MGS12] Vahideh H Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research*, 37(4):559–573, 2012.

[MR13] Marco Molinaro and R Ravi. The geometry of online packing linear programs. *Mathematics of Operations Research*, 39(1):46–59, 2013.

[MSVV07] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22, 2007.

[MY11] Mohammad Mahdian and Qiqi Yan. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing lps. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 597–606. ACM, 2011.

[Ser74] Robert J Serfling. Probability inequalities for the sum in sampling without replacement. *The Annals of Statistics*, pages 39–48, 1974.

[Sho00] Galen R Shorack. *Probability for statisticians*. Springer, 2000.